



**Europäisches
Patentamt**

**European
Patent Office**

**Office européen
des brevets**

J1017 U.S. PTO
09/855863
05/14/01

Bescheinigung

Certificate

Attestation

Die angehefteten Unterla-
gen stimmen mit der
ursprünglich eingereichten
Fassung der auf dem näch-
sten Blatt bezeichneten
europäischen Patentanmel-
dung überein.

The attached documents
are exact copies of the
European patent application
described on the following
page, as originally filed.

Les documents fixés à
cette attestation sont
conformes à la version
initialement déposée de
la demande de brevet
européen spécifiée à la
page suivante.

Patentanmeldung Nr. Patent application No. Demande de brevet n°

00114341.1

Der Präsident des Europäischen Patentamts;
Im Auftrag

For the President of the European Patent Office

Le Président de l'Office européen des brevets
p.o.

I.L.C. HATTEN-HECKMAN

DEN HAAG, DEN
THE HAGUE,
LA HAYE, LE

07/02/01

This Page Blank (uspto)



Europäisches
Patentamt

European
Patent Office

Office européen
des brevets

**Blatt 2 der Bescheinigung
Sheet 2 of the certificate
Page 2 de l'attestation**

Anmeldung Nr.:
Application no.:
Demande n°: 00114341.1

Anmeldetag:
Date of filing:
Date de dépôt: 04/07/00 ✓

Anmelder:
Applicant(s):
Demandeur(s):
International Business Machines Corporation
Armonk, NY 10504
UNITED STATES OF AMERICA

Bezeichnung der Erfindung:
Title of the invention:
Titre de l'invention:

Method and system of weighted context feedback for result improvement in information retrieval

In Anspruch genommene Priorität(en) / Priority(ies) claimed / Priorité(s) revendiquée(s)

Staat:
State:
Pays:

Tag:
Date:
Date:

Aktenzeichen:
File no.
Numéro de dépôt:

Internationale Patentklassifikation:
International Patent classification:
Classification internationale des brevets:

/

Am Anmeldetag benannte Vertragsstaaten:
Contracting states designated at date of filing: AT/BE/CH/CY/DE/DK/ES/FI/FR/GB/GR/IE/IT/LI/LU/MC/NL/PT/SE/TR ✓
Etats contractants désignés lors du dépôt:

Bemerkungen:
Remarks:
Remarques:

This Page Blank (uspto)

- 1 -
D E S C R I P T I O NEPO - Munich
41
04. Juli 2000**Method and System of Weighted Context Feedback for Result
Improvement in Information Retrieval**

BACKGROUND OF THE INVENTION

The present invention relates to the field of document based information retrieval in the intranet and the Internet domain, and in particular to a method for re-ranking an existing result list of documents.

Nowadays digital information systems provide easy access to large amounts of information. For example, users can access great quantities of information in databases on a network or even in personal computers. Mere access to large amounts of information has only limited value, however, without knowing what is useful about the information.

Searching for a certain information in a huge amount of almost unstructured information can be a very cumbersome task. As well as the probability is relatively high that the desired information is somewhere existent in an existing collection of information and potentially can be found, it is at the same time covered up with all that additional unwanted information.

To retrieve a desired piece of information there have been proposed several methods. A first approach is to structure raw data before a search is started. Structuring the raw data in advance can be done by introducing a set of taxonomies (like "cars", "health", "computers" etc.) and assigning each retrieved document into one or more of those categories what can be performed in advance, e.g. by an administrator.

It is noteworthy that the term "document" is understood in the present context in the broadest sense i.e. including any continuous amount of information.

Before a search is started, the user has to preselect one or more categories and thereby reduces the possible amount of returned information, because only the information stored in the selected categories is returned. A simple way to accomplish this is for example having several different indexes for the search engine. The user can then select either one or more indexes as a basis for the search.

The drawback of the above described approach is that extra effort is necessary for preparing the raw data by defining the taxonomy set and by assigning each document into one or more categories (or indexes). Since the information most often is of a dynamic nature, an update of new information or categories (or both) has to be done on a regular basis. Further there is a certain chance that some information is lost because of a wrong assignment to categories or a missing selection by the user.

A second approach is to structure a search result after having finished the search. Structuring the result of a search is either based on the input the user made when he started the search with some query terms, or an attempt is made to dynamically find similarities inside of the documents and group (cluster) them together.

The second approach can be implemented by way of clustering the results which means finding "categories" dynamically by looking for similarities inside of the returned documents. This can be achieved according to several different criteria, for example by scanning for lexical affinities (anyhow related expressions) and bundling those documents that have a certain similarity. Thereby the total list of returned documents is split into several non-overlapping clusters that contain documents which are assumed to deal with the same or similar context.

The drawback of the above second approach is that the search engine has no idea which context the user is really looking for. Moreover, the clustering of the documents is performed on the basis of word tuples that occur with a certain relation to each other in the documents. Ambiguous terms can cause some documents to be scattered all over the clusters, although from the user point of view they deal with the same context. To find those documents the user has to open and read lots of uninteresting information.

An alternative way to implement the second approach, i.e. structuring the search result after having finished the search, is to sort the found documents in a descending order, derived by some comparable criterion. This method is commonly known as "Ranking". The appearance of the search terms in each of the documents is a measurement for the importance of the individual document. All values are normalized so the higher the rank value is, the more importance is assumed.

Various different algorithms are used to determine the individual rank value, most often the document sizes, the index sizes, the total number of returned information or other criteria are considered.

Still another way to implement the second approach is refining the search by adding more precise search terms known as "Narrow Query". The user starts the search process with an initial query and examines some of the returned documents. For each document he assigns a relevance value that reflects if the appropriate document is of high or low value. The search engine scans the marked documents for terms that occur with a high frequency and uses those terms to synthesize a new query. This new query favors the high frequent terms of the documents with a high relevance and excludes the terms of the documents with a low relevance.

But its drawback is that for good results the user has to examine and mark a lot of documents, otherwise the refined query is more or less random.

The predescribed approaches have common the drawback that only the entered search terms can be taken into account. Since search terms often are ambiguous, they can occur in various totally different context and cause a lot of unwanted information to be returned. If, in addition, there are only a few terms entered, there is a high probability that lots of documents get the same high rank value.

Further, in an article by R. Fagin and E. L. Wimmers entitled "Incorporating user preferences in multimedia queries", and published in Proc. 1997 International Conference on Database Theory, pp. 247-261, according to a third approach, it is proposed to weight the search terms in information retrieval with the help of user feedback and to allow to apply weights to any sort of rules. In particular, a formula is described to re-sort and re-rank the result list of a query according to the weighting of the original search terms by the user.

The drawback of the above third approach is that from common experiences in information retrieval it is known that the typical query consists of very few search terms (< 2 terms). In the vast majority of all searches trying to apply any weighting to them would have very little impact. The main reason for the fact that the average search is made with only one or two terms is the lacking of conception of the user what additional context information he had to enter to improve the query and filter the unwanted documents. Following this scenario the user would have to read some of the returned documents first to obtain some context information which he in turn could then use to create additional query terms. Only with a minimum set of query terms that reflect the desired context this method would make sense.

Finally, all of the above cited prior art approaches have common the drawback that, if the preparation of the raw data as described above is not feasible or in a situation, where information can be lost because of a user error, e.g. by selecting a wrong index, only a method that does a post search improvement is acceptable.

Thereupon, all of the described approaches require a relative high effort with opening and reading the contents of documents, thereby wasting time with useless information.

SUMMARY OF THE INVENTION

The object of the present invention therefore is to provide an information retrieval method and system which presents only the desired information without wasting a user's time with opening and reading documents of no value.

Another object is to provide a method and a system for ranking information comprised of a list of documents which presents the desired information on top of a ranking hierarchy.

Still another object is to minimize the possibility of user errors, in particular with respect to generating search queries.

These problems are solved by the features of the independent claims. Advantageous embodiments of the invention are subject matter of the dependent claims.

The invention accomplishes the foregoing by gathering context information from the documents, generating at least one rank criterion from the context information and re-ranking the documents, based on the at least one rank criterion.

The concept underlying the invention is a post search improvement, particularly a human feedback to improve search

results. It is emphasized hereby that the proposed re-ranking is only based on the results of a first search procedure.

The criterion for re-ranking hereby consists of terms themselves found in the documents. These terms contain information which allow to feedback to the context of the underlying documents. These terms, in the following designated as "context terms", are used in a weighting formula in order to separate documents with a matching context from documents with a mismatching context.

The ambiguity of search terms which causes unwanted information to be returned, in a preferred embodiment, is solved by means of context information which is used to filter documents with unwanted context. In another embodiment a context weighting can be performed on a client data processing system, whereby no iterative and time consuming access to a search engine is necessary.

Using lexical affinities to be evaluated as a criterion for re-ranking, is only exemplarily and can be accomplished also by other means that are able to provide a notion of context. A suitable approach could be to use 'features' that have been extracted from the documents or to simply create 'statistics of words'. 'Features' are names of people, places, companies etc. and can be presented as a context indication very similar to lexical affinities. Especially when several extracted features are grouped and weighted together, the context definition would be very precise. 'Word statistics' describe the frequency a given word is occurring within a given document.

In case one finds several different words in a document at least a minimum amount of times - no matter if they occur with a certain proximity to each other like with lexical affinities - and the same is true for another document, it is assumed that both documents are very close to each other concerning their context.

Both information (features as well as word statistics) can also be created by the search engine and transferred to the client together with the result list.

Advantageously, the number of unnecessarily opened documents is hereby essentially reduced since an appropriate weighting of the context terms allows to precise the right context for the desired information.

BRIEF DESCRIPTION OF THE DRAWINGS

The invention will be understood more readily from the following detailed description when taken in conjunction with the accompanying drawings, in which

- Fig. 1 is a combined block and flow diagram depicting a Search process with weighted context feedback by Lexical affinities, in accordance with the invention;
- Fig. 2 is a screenshot of a possible implementation of this method; and
- Fig. 3 is a flow diagram depicting an embodiment of the re-ranking mechanism proposed by the invention.

DETAILED DESCRIPTION OF THE DRAWINGS

Referring to Fig. 1, a combined block and flow diagram is depicted which shows a typical search with weighted context feedback by lexical affinities. The scenario described in the following assumes that it is accepted to start a search with rather unprecise search terms and to try to structure the huge amount of returned data according to context information.

A user 100 starts a search by entering search term(s). The search term(s) is (are) transferred to a search engine 110 which generates a result list 120 ranked by the search term(s). The search engine 110, in parallel, automatically retrieves context information 130 from returned result list 120 which is related 140 to the original list of documents. The context information, in a preferred embodiment, is based on lexical affinities which will be discussed in more detail later.

The search engine 110 presents 150 the context information 130 to the user 100 and asks for a feedback. The user 100 performs a weighting 160 of the presented context information 130 in a range from 'important' to 'non-important'.

The result list 120 is then re-ranked 170 with the user-weighted context information 180 to increase the "rank distance" of important and non important documents. The re-ranking is performed by utilizing a context ranking formula which is described in more detail below. The documents that are now on top of the list (highest context-weighted ranking value) represent the desired information.

Optionally, the re-ranking can be performed iteratively whereby the modified list of documents ranked by the user-weighted context information 180 is created new 190 with different weightings, based on the original list, and re-ranked again performing the predescribed steps.

The underlying re-ranking algorithm is based on an adaptation of the prementioned Fagin's and Wimmers' formula. The way to generate context information is the extraction of "lexical affinities". This method produces pairs of terms that are found in a certain relation within the documents. If one, for example, is searching for the term "display", that term could occur together with the terms "window" or "show" or "computer" or various other combinations. Each combination is an indication for the context the documents are dealing with.

In the formula of R. Fagin and E. L. Wimmers, the following definitions are assumed in the following.

Let x_1, \dots, x_m be the terms that represent the context information. The function $f(x_1, \dots, x_m)$ is then the ranking function that returns a rank value for a given document with respect to these terms x_1, \dots, x_m . $\theta = (\theta_1, \dots, \theta_m)$ is the overall weighting for all of the m terms and θ_i is the individual weight for term x_i . Under the assumption that $\theta_1, \dots, \theta_m$ are all non-negative, $\theta_1 \geq \dots \geq \theta_m$ and they all sum to one, where we use the weighted ranking formula according to Fagin and Wimmers which is:

$$(\theta_1 - \theta_2)f(x_1) + 2(\theta_2 - \theta_3)f(x_1, x_2) + 3(\theta_3 - \theta_4)f(x_1, x_2, x_3) + \dots + m\theta_m f(x_1, \dots, x_m) \quad (1)$$

A scenario where all m context terms get a different rating would increase the necessary calculating effort to the theoretical maximum of m . According to equation (1) we had to calculate the ranking value for each document m times, i.e.

$$f(x_1), f(x_1, x_2), f(x_1, x_2, x_3), \dots, f(x_1, \dots, x_m),$$

since the expressions $(\theta_n - \theta_{n+1})$ in equation (1) would all become non-equal to zero. Besides the enormous visualization effort to allow a user to choose m different ratings for all the m context terms, it seems very unlikely that anybody would really do this. An approach where for example three different rating levels are selectable is much more practical.

Let the rating levels be 'high', 'medium' and 'low' and the relation between those three rating levels

$$\theta_a = 2 \theta_b \text{ and } \theta_c = 0.5 \theta_b \quad (2)$$

During further evaluation with more samples of test documents it might turn out, that a stronger spreading of those three levels is more useful, such as $\theta_a = 4\theta_b$ and $\theta_c = 0.25 \theta_b$ or even more.

Now the rating selections of all of the context terms made by the user can be divided into

a terms rated with 'high'

b terms rated with 'medium'

c terms rated with 'low' (3)

where all selections sum up to the total number of context terms :

$$a + b + c = m \quad (4)$$

Because of the assumption that all $\theta_1, \dots, \theta_m$ are non-negative and sum to one, it can be written

$$a \theta_a + b \theta_b + c \theta_c = 1 \text{ and}$$

$$\theta_b = (1 - a \theta_a - c \theta_c) / b.$$

Because of (2) it follows

$$\theta_b = (1 - 2a \theta_b - 0.5c \theta_b) / b$$

and finally

$$\theta_b = 1 / (2a + b + 0.5c) \quad (5)$$

Allowing only three discrete ranking levels, one gets only three groups of θ_n sorted according to their importance to the user. The first 'a' context terms are rated high, the next 'b' context terms are rated medium and the last 'c' context terms are rated low.

$$\begin{aligned}\theta_1 &= \dots = \theta_a \text{ and} \\ \theta_{a+1} &= \dots = \theta_{a+b} \text{ and} \\ \theta_{a+b+1} &= \dots = \theta_m\end{aligned}\tag{6}$$

Using equation (6), equation (1) can be simplified into

$$\begin{aligned}& a(\theta_a - \theta_{a+1}) \cdot f(x_1, x_2, \dots, x_a) \\ & + (a+b)(\theta_{a+b} - \theta_{a+b+1}) f(x_1, x_2, \dots, x_{a+b}) \\ & + (a+b+c) \theta_c f(x_1, x_2, \dots, x_{a+b+c})\end{aligned}\tag{7}.$$

In case of the only three relevance values 'a', 'b', and 'c', $\theta_{a+1} = \theta_b$, $\theta_{a+b} = \theta_b$ and $\theta_{a+b+1} = \theta_c$, one can substitute

$$\begin{aligned}& a(\theta_a - \theta_b) f(x_1, x_2, \dots, x_a) \\ & + (a+b)(\theta_b - \theta_c) f(x_1, x_2, \dots, x_{a+b}) \\ & + (a+b+c) \theta_c f(x_1, x_2, \dots, x_{a+b+c})\end{aligned}\tag{8}.$$

This has to be converted into a normalized form with only dependencies to the number of context terms rated with the three discrete rating levels and the rank functions of the appropriate context terms.

Following equation (2) one can substitute equation (8) with

$$\begin{aligned}& a \theta_b f(x_1, \dots, x_a) \\ & + 0.5(a+b) \theta_b f(x_1, \dots, x_{a+b}) \\ & + 0.5(a+b+c) \theta_b f(x_1, \dots, x_{a+b+c})\end{aligned}\tag{9}$$

and according to equation (5) one finally obtains

$$\begin{aligned} & [2a f(x_1, \dots, x_a) \\ & + (a+b) f(x_1, \dots, x_{a+b}) \\ & + (a+b+c) f(x_1, \dots, x_{a+b+c})] / (4a+2b+c) \end{aligned} \quad (10)$$

Equation (10) calculates the relevance of a document with respect to the context terms x_1, \dots, x_m when a , b and c are the number of terms that have been assigned a high (a), medium (b) and low (c) relevance and $f(x_1, \dots, x_a)$, $f(x_1, \dots, x_{a+b})$ and $f(x_1, \dots, x_{a+b+c})$ are the partial relevance functions of this document with respect to a subset of context terms.

In comparison to equation (1), one now has to call the ranking function only three times with different parameter lists which saves a considerable amount of processing effort. The additional factors in equation (10) are necessary to fulfill the three desiderata of equation (1). For the further details concerning these factors it is referred to the prementioned article by Fagin and Wimmer.

A first way is to feed back the weighting information to the Text Search Engine and calculate the new rank value for all documents of the result list. This reveals some major problems as it is currently not possible to call a ranking function with new parameters (the context terms) on the basis of a prior search, because the statistics of these new parameters are not available. Only the original search terms, that are known during the search process, are complemented by the statistics information of the search indexes and could then be used for a ranking.

Besides this problem, it would also require some additional communication with the search engine that increases the response time to the user.

Therefore it is proposed a new and very efficient ranking algorithm to be used solely on the client side without any further access to the Text Search Engine. The processing effort is relatively low, so that it can also be used on so-called "thin clients" and a quick response to the user is guaranteed.

The new ranking algorithm is based on the "absolute" ranking of the prior search that was derived from the query terms only. Depending on the occurrence of the context terms in a given document and their importance for the user a new rank value is calculated:

Let R_d be the "absolute" rank value of a given document 'd' that resulted from the search, $T_d = (x_1, \dots, x_n)$ is the tuple of context terms that are contained in document 'd', then is

$$f_d(x_1, \dots, x_n) = R_d \quad \text{if } x_1, \dots, x_n \text{ element of } T_d$$

and

$$f_d(x_1, \dots, x_n) = 0 \quad \text{if } x_1, \dots, x_n \text{ non-element of } T_d \quad (11)$$

This new ranking equation (11) is used together with the adaptation of Fagin's and Wimmer's weighted ranking (10) and results in a new rank value for each document. The desired effect is a change in the result list order. Those documents that contain the preferred context terms will appear on top of the list since their new rank value is higher than before. The documents that contain only the context terms that got a low relevance from the user will be pushed to the end of the list since their rank value is lowered.

The above described (re-)ranking algorithm can be implemented in a Java environment with the help of Visual Builders. For the special purpose of weighted context feedback a new Java Bean was developed and seamlessly integrated into the Java Beans component model architecture.

Fig. 2 depicts a screenshot of a possible implementation of this method. The right side of this screenshot shows the result of a search in a table format. According to the ranking of the result list by the query term (in this case it was the single word "display"), the documents are sorted top-down.

On the left side there is the visualization of the Context Feedback. The left column with the header title "Importance" holds three graphical buttons on each line, the right column with the header title "Context" holds the pairs of lexical affinities on each line. The table is scrollable up and downward to display all the context information and the corresponding weighting.

Immediately after having performed a search, the right column is filled with all context information that could be extracted from all documents of the result list, the weighting of the individual context information in the left column is all set to medium.

By clicking the plus or minus symbols the weighting for each pair of lexical affinities can be modified. Leaving some entries in the "medium" state is also possible. As soon as the user is done with the weighting, he presses the "Rank new" button and the information is sent to the search client to do a weighted re-rank. This causes the right result list table to refresh its contents with the new ranking order.

Fig. 3 shows the mechanism that leads to a new ranking value 'Rd_new' for each document 'd' contained in a result list. As already mentioned, the number of iterations per individual document depends on the number of discrete weighting levels that are offered to the user. Following the example depicted in Fig. 2 with three different ratings ('high', 'medium' and 'low') one has to calculate up to three portions per document that make up the new rank value.

For each document it is started with its 'absolute' rank value of the unweighted result list 'Rd' and set 300 the initial value of 'Rd_new' equal to '0'. Further it is checked 310 whether any of the lexical affinities that have been weighted 'high' are contained in document 'd'. If true 320, the new ranking value 'Rd_new' is set 330 to ' $2*a*Rd$ ' ('a' is the number of lexical affinities weighted with 'high'), otherwise 'Rd_new' remains unchanged 340.

In a next iteration 350 one performs a similar calculation for all the lexical affinities that have been rated 'high' and 'medium', and finally a third iteration 360 takes into account all the 'high', 'medium' and 'low' rated lexical affinities. The last step is a normalization of the new rank value 'Rd_new'.

It is noted hereby that the above described mechanism for re-ranking documents can be applied to any text documents or non-text documents, i.e. all kinds of documents where a ranking principally can be performed. The term "documents" is understood in its broadest sense and means an arbitrary allied amount of information.

This Page Blank (uspto)

C L A I M S

1. A method for ranking a list of documents, in particular for re-ranking an existing result list of documents, comprising the steps of

gathering context information from the documents;

generating at least one rank criterion from the context information;

ranking the documents, based on the at least one rank criterion.
2. Method according to claim 1, whereby utilizing lexical affinities as the context information.
3. Method according to claim 1 or 2, whereby utilizing features as the context information.
4. Method according to any of claims 1 to 3, whereby utilizing word frequency statistics as the context information.
5. Method according to claim 2, whereby extracting lexical affinities from the documents in order to obtain the context information.
6. Method according to claim 3, whereby extracting features from the documents in order to obtain the context information
7. Method according to claim 4, whereby extracting word frequency statistics from the documents in order to obtain the context information

8. Method according to any of claims 1 to 7, comprising the further step of weighting of the context information by a weighting function.
9. Method according to claim 8, whereby providing discrete ranking levels.
10. Method according to any of the preceding claims for re-ranking an existing list of text documents, comprising the steps of

detecting of lexical affinity terms contained in the documents;

presenting the lexical affinity terms to a user;

gathering user preferences for the lexical affinity terms;

re-ranking the documents based on the user preferences.
11. Method according to any of the preceding claims for re-ranking an existing list of text documents, comprising the steps of

detecting of feature terms contained in the documents;

presenting the feature terms to a user;

gathering user preferences for the feature terms;

re-ranking the documents based on the user preferences.

12. Method according to any of the preceding claims for re-ranking an existing list of text documents, comprising the steps of
- creating of word frequency statistics from the documents;
- presenting the words with a minimum frequency to a user;
- gathering user preferences for the presented words of a minimum frequency;
- re-ranking the documents based on the user preferences.
13. Method according to any of the preceding claims, whereby the re-ranking is based on the original ranking position of the documents.
14. Method according to any of the preceding claims, whereby using ranking equation (eq. (11)) together with a weighted ranking equation (eq. (10)).
15. System for ranking a list of documents, in particular for re-ranking an existing result list of documents, Comprising means for gathering context information from the documents, for generating at least one rank criterion from the context information and for re-ranking the documents, based on the at least one rank criterion.
16. System according to claim 15, comprising means for extracting lexical affinities from the documents in order to obtain the context information.
17. System according to claim 15 or 16, comprising means

for weighting of the context information by a weighting function.

04. Juli 2000

A B S T R A C T

Disclosed is a method and a system for re-ranking an existing result list of documents. A user (100) starts a search by entering search term(s). The search term(s) is (are) transferred to a search engine (110) which generates a result list (120) ranked by the search term(s). The search engine (110), in parallel, automatically retrieves context information (130) from returned result list (120) which is related (140) to the original list of documents. The search engine (110) presents (150) the context information (130) to the user (100) and asks for a feedback. The user (100) performs a weighting (160) of the presented context information (130) in a range from 'important' to 'non-important'. The result list (120) is then re-ranked (170) with the user-weighted context information (180) to increase the "rank distance" of important and non important documents. The documents that are now on top of the list (highest context-weighted ranking value) represent the desired information. The underlying re-ranking algorithm is based on an adaptation of a formula by Fagin and Wimmers. The way to generate context information is the extraction of "lexical affinities". This method produces pairs of terms that are found in a certain relation within the documents.

(Fig. 1)

This Page Blank (uspto)

(Drawings)


EPO - Munich
41
04. Juli 2000

FIG. 1

Importance	Content	Rank	Document ID	Collection	Size
1	beam,electron	40	pat2426810	SCIENCE	313
2	beam,ray	39	pat2426500	SCIENCE	273
3	film,measure	34	pat2407474	SCIENCE	187
4	data,display	32	pat2504914	SCIENCE	2420
5		32	pat2492304	SCIENCE	216
6		30	pat2453263	SCIENCE	262
7		30	pat2426160	SCIENCE	318

FIG. 2

